



HAL
open science

Risk Bounds for CART Classifiers under a Margin Condition

Servane Gey

► **To cite this version:**

| Servane Gey. Risk Bounds for CART Classifiers under a Margin Condition. 2009. hal-00362281v4

HAL Id: hal-00362281

<https://hal.science/hal-00362281v4>

Preprint submitted on 2 Jul 2010 (v4), last revised 1 Mar 2012 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Risk Bounds for CART Classifiers under a Margin Condition

Servane Gey*

Abstract

Risk bounds for Classification And Regression Trees (CART) classifiers are obtained under a margin condition in the binary supervised classification framework. These risk bounds are derived conditionally on the construction of the maximal binary tree and permit to prove that the linear penalty used in the CART pruning algorithm is valid under a margin condition. It is also shown that, conditionally on the construction of the maximal tree, the final selection by test sample does not alter dramatically the estimation accuracy of the Bayes classifier.

In the two-class classification framework, the risk bounds obtained by using penalized model selection validate the CART algorithm which is used in many data mining applications in Biology, Medicine or Image Coding for instance.

Keywords: Classification, CART, Pruning, Margin, Risk Bounds.

AMS 2000 subject classifications : primary 62G99, 62H30 and secondary 62-07.

1 Introduction

The Classification And Regression Trees (CART) method proposed by Breiman, Friedman, Olshen and Stone [9] in 1984 consists in constructing an efficient algorithm that gives a piecewise constant estimator of a classifier or a regression function from a training sample of observations. This algorithm is based on binary tree-structured partitions and on a penalized criterion that selects “good” tree-structured estimators among a huge collection of trees. It currently yields some easy-to-interpret and easy-to-compute estimators which are widely used in many applications in Medicine, Meteorology, Biology, Pollution or Image Coding (see [10], [39] for example). This kind of algorithm is often performed when the space of explanatory variables is high-dimensional. Due to its recursive computation, CART needs few computations to provide convenient classifiers, which accelerates the computation time drastically when the number of variables is large. It is now widely used in the genetics framework (see [16] for example), or more generally to reduce variable dimension (see [33] [26] for example).

The CART algorithm provides classifiers or regressors represented by binary decision trees. An example of the latter is given in Figure 1. Suppose we have a couple of covariates (X_1, X_2) belonging to $[0; 1]^2$. The partition is defined recursively by a sequence of questions asked at each node of the tree: if the answer is positive, go to the left node, if not, go to the right node. Hence the first question corresponds to a two-part partition of the covariate space. Then, each part is split into two subparts, and so on. Hence each node of the tree represents a subset of the covariates space defined by the successive questions.

*Laboratoire MAP5 - UMR 8145, Université Paris Descartes, 75270 Paris Cedex 06, France. Servane.Gey@parisdescartes.fr

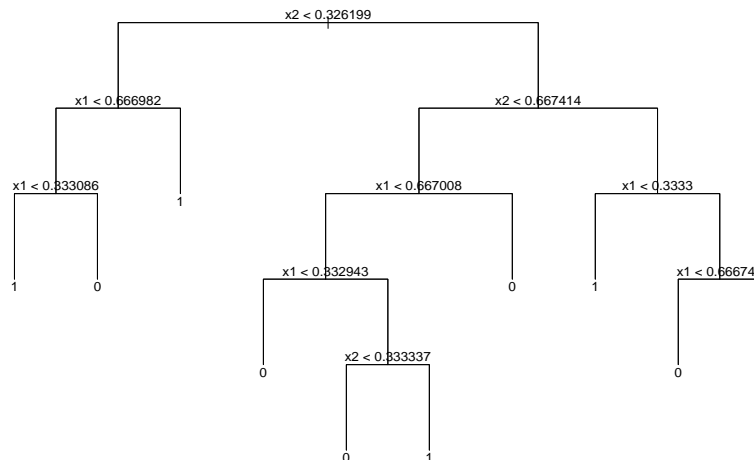


Figure 1: Decision tree example

The final partition is given by the leaves of the tree. Finally, a predictive value for the dependent variable is associated to each leaf.

To construct such a tree from a training sample of observations, the CART algorithm consists in constructing a large dyadic recursive tree from the observations by minimizing some local impurity function at each step. Then, the constructed tree is pruned to obtain a finite sequence of nested trees thanks to a penalized criterion, whose penalty term is proportional to the number of leaves. The linearity of the penalty term is fundamental to ensure that the whole information is kept in the obtained sequence. CART differs from the algorithm proposed by Blanchard *et al.* [5] by the fact that the first large tree is constructed locally, and not in a global way by minimizing some loss function on the whole sample. For further results on the construction of the large tree, we refer to Nobel [30, 31], and Nobel and Olshen [32] about Recursive Partitioning.

In this paper, our concern is the pruning step which entails the choice of the penalty function. Gey *et al.* [17] gave an answer to this question in the regression framework. Following this previous work, the present paper aims at validating the choice of the penalty in the two class classification framework. In what follows, we establish the link between the CART algorithm and a model selection procedure, where the collection of models is a collection of random decision trees constructed on the training sample of observations. In its pruning procedure, CART selects a small collection of trees within the whole collection of random trees. Then, a final tree belonging to the small collection is selected either by cross-validation or by test sample. The present paper focuses on the test sample method. We exhibit risk bounds for the chosen tree under some conditions on the joint distribution of the variables. These risk bounds validate the choice of the penalty used in the pruning step, and show that the impact of the selection via test sample is conveniently controlled.

The CART method takes place in the following general classification framework. Suppose one observes a sample \mathcal{L} of N independent copies $(X_1, Y_1), \dots, (X_N, Y_N)$ of the random variable (X, Y) , where the explanatory variable X takes values in a measurable space

\mathcal{X} and is associated with a label Y taking values in $\{0, 1\}$. A classifier is then any function f mapping \mathcal{X} into $\{0, 1\}$. Its quality is measured by its misclassification rate

$$P(f(X) \neq Y),$$

where P denotes the joint distribution of (X, Y) . If P were known, the problem of finding an optimal classifier minimizing the misclassification rate would be easily solved by considering the Bayes classifier f^* defined for every $x \in \mathcal{X}$ by

$$f^*(x) = \mathbb{1}_{\eta(x) \geq 1/2}, \quad (1)$$

where $\eta(x)$ is the conditional expectation of Y given $X = x$, that is

$$\eta(x) = P[Y = 1 \mid X = x]. \quad (2)$$

As P is unknown, the goal is to construct from the sample $\mathcal{L} = \{(X_1, Y_1), \dots, (X_N, Y_N)\}$ a classifier \tilde{f} that is as close as possible to f^* in the following sense: since f^* minimizes the misclassification rate, \tilde{f} will be chosen in such a way that its misclassification rate is as close as possible to the misclassification rate of f^* , i.e. in such a way that the expected loss

$$l(f^*, \tilde{f}) = P(\tilde{f}(X) \neq Y) - P(f^*(X) \neq Y) \quad (3)$$

is as small as possible. Then, the quality of \tilde{f} will be measured by its risk, i.e. the expectation with respect to the \mathcal{L} -sample distribution

$$R(\tilde{f}, f^*) = \mathbb{E}[l(f^*, \tilde{f})]. \quad (4)$$

Numerous works have dealt with the issue of predicting a label from an input $x \in \mathcal{X}$ via the construction of a classifier (see for example [1], [38], [11], [34], [18]). There is a large collection of methods coming both from computational and statistical areas and based on learning a classifier from a learning sample, where the inputs and labels are known. For a non exhaustive yet extensive bibliography on this subject, we refer to Boucheron *et al.* [6]. We based our computation of risk bounds for the CART classifier on recent results (see for instance [25], [35], [36], [29], [20, 21], [28], [22], [19]). They stem from Vapnik's results (see [37], [23] for example), which show that, without any assumption on the joint distribution P , the penalty term used in the model selection procedure is proportional to the square root of the number of leaves over N . Nevertheless, it has also been shown that, under the overoptimistic zero-error assumption (that is $Y = \eta(X)$ almost surely, where η is defined by (2)), this penalty term is proportional to the number of leaves over N .

In fact, these two extreme cases can be modulated by so-called *margin assumptions*, which permit to compare the loss of a classifier with its \mathbb{L}^2 distance to the Bayes classifier f^* . Numerous margin assumptions have been investigated by the above-cited authors; some permit to obtain penalty terms proportional to the number of leaves over N to the power κ , with $1/2 \leq \kappa \leq 1$ (see for example [36] and [29]). Hence these margin assumptions make a link between the "global" pessimistic case (without any assumption on P) and the zero-error case. More recent works (see [20, 21], [2] for instance) deal with data-driven penalties based on local Rademacher complexities and use more general margin assumptions than those proposed in [25] and [29]. Those works also show that the margin assumption necessary to obtain a penalty term proportional to the number of leaves over N is one of the strongest. Let us introduce the following margin assumption:

$$\mathbf{MA}(1) \exists h \in]0; 1[\quad \forall f : \mathcal{X} \rightarrow \{0; 1\} \quad l(f^*, f) \geq h \mathbb{E} [(f(X) - f^*(X))^2],$$

where l is the expected loss defined by (3). Margin assumption **MA(1)** is implied by the more intuitive assumption proposed by Massart *et al.* in [29] (see also the slightly weaker condition proposed in [19]):

$$\mathbf{MA(2)} \quad \exists h \in]0; 1[\quad \mathbb{P}(|2\eta(X) - 1| \leq h) = 0.$$

Assumption **MA(2)** means that (X, Y) is sufficiently well distributed to ensure that there is no region in \mathcal{X} for which the toss-up strategy could be favored over others: h can be viewed as a measurement of the gap between labels 0 and 1 in the sense that, if $\eta(x)$ is too close to 1/2, then choosing 0 or 1 will not make a real difference for that x .

Below, we prove that, under **MA(1)**, the penalty used by CART in the pruning step is convenient.

In the rest of the paper, the constant h will denote the so-called *margin*. Of course margin assumption **MA(1)** is chosen for its relevance in the particular framework of CART and shall be adapted, or simply ignored, depending on the problem under study.

As mentioned above, we leave aside the construction of the first large tree. Thus, all our upper bounds for the risk of the classifier obtained by CART are considered conditionally on the recursive construction of the first large tree, called *maximal tree*. Moreover, we focus on non-asymptotic bounds.

We also leave aside the problem of consistency of CART. CART is known to be nonconsistent in many cases. Some results and conditions to obtain consistency can be found in the paper by Devroye *et al.* [11]. Section 3 briefly presents consistent results for CART based on the risk bounds obtained.

We focus on two methods that use a test sample: let us split \mathcal{L} in three independent subsamples \mathcal{L}_1 , \mathcal{L}_2 and \mathcal{L}_3 , containing respectively n_1 , n_2 and n_3 observations, with $n_1 + n_2 + n_3 = N$. \mathcal{L}_1 , \mathcal{L}_2 and \mathcal{L}_3 are taken randomly in \mathcal{L} , except if the design is fixed. In that case one takes, for example, one observation out of three to obtain each subsample. Given these three subsamples, suppose that either a large tree is constructed using \mathcal{L}_1 and then pruned using \mathcal{L}_2 (as done in Gelfand *et al.* [14]), or a large tree is constructed and pruned using the subsample $\mathcal{L}_1 \cup \mathcal{L}_2$ (as done in [9]).

Then the final step used in both cases is to choose a subtree among the sequence by making \mathcal{L}_3 go down each tree of the sequence and selecting the tree having the minimum empirical misclassification rate : for $k = 1, 2, 3$ and for f a classifier, the empirical misclassification rate of f on \mathcal{L}_k is given by:

$$\gamma_{n_k}(f) = \frac{1}{n_k} \sum_{(X_i, Y_i) \in \mathcal{L}_k} \mathbb{1}_{Y_i \neq f(X_i)}. \quad (5)$$

The final estimator \tilde{f} of f^* is defined by:

$$\tilde{f} = \operatorname{argmin}_{\{1 \leq i \leq K\}} \left[\gamma_{n_3}(\hat{f}_{T_i}) \right], \quad (6)$$

where \hat{f}_{T_i} is the piecewise binary estimator of f^* defined on the leaves of the tree T_i and K is the number of trees appearing in the sequence.

The paper is organized as follows. Section 2 gives an overview of the CART algorithm, and introduces the methods and notations used in the following sections. Section 3 presents the main theoretical results for classification trees: Theorem 1 bears on the whole algorithm, while Propositions 1, 2 concern the pruning procedure and Proposition 3 concerns the final step. Section 4 offers prospects about the margin effect on classification trees. Proofs are gathered in Section 5.

2 The CART Procedure

Let us give a short account of the CART procedure in the classification case and recall the results associated with it, which are fully explained in [9].

CART is based on a recursive partitioning using a training sample $\tilde{\mathcal{L}}$ of the random variable $(X, Y) \in \mathcal{X} \times \{0, 1\}$ ($\tilde{\mathcal{L}} = \mathcal{L}_1$ or $\tilde{\mathcal{L}} = \mathcal{L}_1 \cup \mathcal{L}_2$), and a class \mathcal{S} of subsets of \mathcal{X} which tells us how to split at each step. For instance, if $\mathcal{X} = \mathbb{R}^d$, \mathcal{S} is usually taken as some class of half-spaces of \mathcal{X} , for example the half-spaces of \mathcal{X} with frontiers parallel to the axes (see for example [9], [12]). Below, we consider a class \mathcal{S} with finite Vapnik-Chervonenkis dimension, henceforth referred to as VC-dimension (for a complete overview of the VC-dimension see [37]).

The procedure is computed in two steps, called the growing procedure and the pruning procedure. The growing procedure permits to construct a maximal binary tree T_{max} from the data by recursive partitioning, and then the pruning procedure permits to select, among all the subtrees of T_{max} , a sequence that contains the entire statistical information.

2.1 Growing and pruning procedures

2.1.1 Growing Procedure

Since our main interest in this paper is the pruning procedure, we present an overview of the growing procedure (for more details about the growing procedure, see [9]).

The growing procedure is based on a recursive binary partitioning of \mathcal{X} . Let us start with the first step: \mathcal{X} is split into two parts by minimizing some empirical convex function on \mathcal{S} . A strictly convex function is used in order to avoid ties, which is systematically the case when using the simplest empirical misclassification rate (see [9], [24]). Thus this function is chosen in such a way that the data are split into two groups where the labels of the data in each group are as similar as possible. It implies that the empirical misclassification rate in each subgroup is largely reduced. Note that the sum of empirical misclassification rates of each subgroup (called *node*) is always smaller than the global empirical misclassification rate on the sample $\tilde{\mathcal{L}}$ (called the *root* t_1 of the tree). In the tree terminology, one adds to the root t_1 a left node t_L and a right node t_R . In what follows, we always assimilate a tree node with its corresponding subset in \mathcal{S} . Finally, a label is given to each node by majority vote (which corresponds to minimizing the empirical misclassification rate in each node).

Then the same elementary step is applied recursively to the two generated subsamples $\{(X_i, Y_i) ; X_i \in t_L\}$ and $\{(X_i, Y_i) ; X_i \in t_R\}$ until some convenient stopping condition is satisfied. This generates the maximal tree T_{max} ; one calls terminal nodes or leaves the final nodes of T_{max} .

2.1.2 Pruning Procedure

Recall that a pruned subtree of T_{max} is defined as any binary subtree of T_{max} having the same root t_1 as T_{max} .

Now, let us introduce some notations:

- (i) Take two trees T_1 and T_2 . Then, if T_1 is a pruned subtree of T_2 , write $T_1 \preceq T_2$.
- (ii) For a tree T , \tilde{T} denotes the set of its leaves and $|\tilde{T}|$ the cardinality of \tilde{T} .

To prune T_{max} , one proceeds as follows. First simply denote by n the number of data used. Notice that, given a tree T and \mathcal{F}_T a set of binary piecewise functions in $\mathbb{L}^2(\mathcal{X})$

defined on the partition given by the leaves of T , one has

$$\begin{aligned}\hat{f}_T &= \operatorname{argmin}_{g \in \mathcal{F}_T} \gamma_n(g) \\ &= \sum_{t \in \tilde{T}} \operatorname{argmax}_{\{Y_i ; X_i \in t\}} |\{Y_i ; X_i \in t\}| \mathbb{1}_t,\end{aligned}$$

where γ_n is the empirical misclassification rate defined by (5) and $\mathbb{1}_t(x) = 1$ if x falls in the leaf t , $\mathbb{1}_t(x) = 0$ otherwise.

Then, given $T \preceq T_{max}$ and $\alpha > 0$, one defines

$$\operatorname{crit}_\alpha(T) = \gamma_n(\hat{f}_T) + \alpha \frac{|\tilde{T}|}{n} \quad (7)$$

the penalized criterion for the so called temperature α , and T_α the subtree of T_{max} satisfying:

- (i) $T_\alpha = \operatorname{argmin}_{T \preceq T_{max}} \operatorname{crit}_\alpha(T)$,
- (ii) if $\operatorname{crit}_\alpha(T) = \operatorname{crit}_\alpha(T_\alpha)$, then $T_\alpha \preceq T$.

Thus T_α is the smallest minimizing subtree for the temperature α . The existence and the unicity of T_α are proved in [9, pp 284-290].

The aim of the pruning procedure is to raise temperature α and to record the corresponding T_α . The algorithm is iterative: it consists in minimizing a function of the nodes at each step, which leads to a finite decreasing sequence of subtrees pruned from T_{max}

$$T_{max} \succeq T_1 \succ \dots \succ T_{K-1} \succ T_K = \{t_1\}$$

corresponding to a finite increasing sequence of temperatures

$$0 = \alpha_1 < \alpha_2 < \dots < \alpha_{K-1} < \alpha_K,$$

where t_1 corresponds to the root of T_{max} as defined in the growing procedure.

Remark 1. T_1 is the smallest subtree for temperature 0, so it is not necessarily equal to T_{max} .

Breiman, Friedman, Olshen and Stone's Theorem [9] justifies this algorithm:

Theorem 2.1.1 (Breiman, Friedman, Olshen, Stone).

The sequence $(\alpha_k)_{1 \leq k \leq K}$ is nondecreasing, the sequence $(T_k)_{1 \leq k \leq K}$ is nonincreasing and, given $k \in \{1, \dots, K\}$, if $\beta \in [\alpha_k, \alpha_{k+1}[$, then $T_\beta = T_{\alpha_k} = T_k$.

This theorem allows us to check that, for any $\alpha > 0$, T_α belongs to the sequence $(T_k)_{1 \leq k \leq K}$. This algorithm significantly reduces the complexity of the choice of a subtree pruned from T_{max} , since by Theorem 2.1.1 the sequence of pruned subtrees contains the whole statistical information according to the choice of the penalty function used in (7). Consequently it is useless to look at all the subtrees. Notice that the form of the penalized criterion is essential to obtain Theorem 2.1.1. Hence, to fully validate this algorithm completely, we need to show that the choice of penalty is relevant.

The final step is to choose a suitable temperature α . Instead of minimizing over α , this issue is dealt with by using a test-sample to provide the final estimator \tilde{f} , as mentioned

in the Introduction, via equality (6). The results given in Sections 3.1 and 3.2 deal with the performance of the piecewise constant estimators given by T_α for α fixed and with the performance of \tilde{f} respectively.

Before focusing on risk bounds, let us present the methods and notations used to obtain these bounds.

2.2 Methods and Notations

For a given tree T , \mathcal{F}_T will denote the set of classifiers defined on the partition given by the leaves of T , that is

$$\mathcal{F}_T = \left\{ \sum_{t \in \tilde{T}} a_t \mathbb{1}_t ; (a_t) \in \{0, 1\}^{|\tilde{T}|} \right\}, \quad (8)$$

where \tilde{T} refers the set of the leaves of T . Thus \hat{f}_T is the empirical risk minimizer classifier on \mathcal{F}_T . For any tree-structured estimator \hat{f} of f^* , \hat{f} is said to satisfy an oracle inequality if there exists some nonnegative constant C , such that

$$R_{\mathcal{L}_1}(\tilde{f}, f^*) \leq C \inf_{T \leq T_{max}} R_{\mathcal{L}_1}(\hat{f}_T, f^*),$$

where, $R_{\mathcal{L}_1}(\cdot, f^*) = \mathbb{E}[l(f^*, \cdot) | \mathcal{L}_1]$, and $\mathbb{E}[\cdot | \mathcal{L}_1]$ denotes the conditional expectation given the subsample \mathcal{L}_1 .

To estimate f^* using the CART algorithm and to compare the performance of \tilde{f} with those of each \hat{f}_T , two different methods can be applied:

- M1: \mathcal{L} is split in three independent parts \mathcal{L}_1 , \mathcal{L}_2 and \mathcal{L}_3 containing respectively n_1 , n_2 and n_3 observations, with $n_1 + n_2 + n_3 = N$. Hence T_{max} is constructed using \mathcal{L}_1 , then pruned using \mathcal{L}_2 and finally the best subtree \hat{T} is selected among the sequence of pruned subtrees thanks to \mathcal{L}_3 , and we define $\tilde{f} = \hat{f}_{\hat{T}}$.
- M2: \mathcal{L} is split in two independent parts \mathcal{L}_1 and \mathcal{L}_3 containing respectively n_1 and n_3 observations, with $n_1 + n_3 = N$. Hence T_{max} is constructed and pruned using \mathcal{L}_1 and finally the best subtree \hat{T} is selected among the sequence of pruned subtrees thanks to \mathcal{L}_3 , and we define $\tilde{f} = \hat{f}_{\hat{T}}$.

Note that a penalty is needed in both methods in order to reduce the number of candidate tree-structured models contained in T_{max} . Indeed, if one does not penalize, the number of models to be considered grows exponentially with N (see [9]). So making a selection by using a test sample without penalizing requires visiting all the models. In that case, looking for the best model in the collection of all subtrees pruned from the maximal one becomes explosive. Hence penalizing permits to reduce significantly the number of trees taken into account; it provides a convenient risk for \tilde{f} .

3 Risk Bounds

This section is devoted to the results obtained on the performance of the CART classifiers for both methods M1 and M2. We shall first present a general theorem, then give more precise results on the last two parts of the algorithm, which are the pruning procedure

and the final selection by test sample.

Assume that the following margin assumption is fulfilled: there exists some absolute constant $h \in]0; 1[$ such that, for every classifier f ,

$$\begin{aligned} \mathbf{MA(1a)} \quad l(f^*, f) &\geq h \mathbb{E} [(f(X) - f^*(X))^2] && \text{if } \tilde{f} \text{ is constructed via M1} \\ \mathbf{MA(1b)} \quad l_{n_1}(f^*, f) &\geq h \frac{1}{n_1} \sum_{X_1^{n_1}} (f(X_i) - f^*(X_i))^2 && \text{if } \tilde{f} \text{ is constructed via M2} \end{aligned}$$

where l is the expected loss (3) defined in Section 1, $X_1^{n_1} = \{X_i; (X_i, Y_i) \in \mathcal{L}_1\}$ and $l_{n_1}(f^*, f)$ is the empirical expected loss conditionally on the grid $X_1^{n_1}$ defined by

$$l_{n_1}(f^*, f) = \mathbb{E}_Y \left[\frac{1}{n_1} \sum_{X_1^{n_1}} \mathbb{1}_{f(X_i) \neq Y_i} - \mathbb{1}_{f^*(X_i) \neq Y_i} \right], \quad (9)$$

with \mathbb{E}_Y the expectation with respect to the marginal distribution of Y .

Theorem 1. *Given N independent pairs of variables $((X_i, Y_i))_{1 \leq i \leq N}$ of common distribution \mathbb{P} , with $(X_i, Y_i) \in \mathcal{X} \times \{0, 1\}$, let us consider the estimator \tilde{f} (6) of the Bayes classifier f^* (1) obtained via the CART procedure as defined in section 2. Then we have the following results.*

(i) *if \tilde{f} is constructed via M1:*

Let $l(f^, \tilde{f})$ be the expected loss (3) of \tilde{f} and h be the margin given by **MA(1a)**. Then there exist some absolute constants C, C_1 and C_2 such that*

$$\mathbb{E} [l(f^*, \tilde{f}) \mid \mathcal{L}_1] \leq C \inf_{T \leq T_{max}} \left\{ \inf_{f \in \mathcal{F}_T} \mathbb{E} [l(f^*, f) \mid \mathcal{L}_1] + h^{-1} \frac{|\tilde{T}|}{n_2} \right\} + h^{-1} \frac{C_1}{n_2} \quad (10)$$

$$+ h^{-1} C_2 \frac{\log(n_1)}{n_3}. \quad (11)$$

(ii) *if \tilde{f} is constructed via M2:*

Let $P_{\mathcal{L}_1}$ be the product distribution on \mathcal{L}_1 , let $l_{n_1}(f^, \tilde{f})$ be the empirical expected loss of \tilde{f} conditionally on the grid $X_1^{n_1}$ (9), and h be the margin given by **MA(1b)**. Let V be the Vapnik-Chervonenkis dimension of the set of splits used to construct T_{max} and suppose that $n_1 \geq V$. Let K be the number of pruned subtrees of the sequence provided by the pruning procedure. Then there exist some absolute constants C', C'_1, C''_1 and C_2 such that, for every $\delta \in]0; 1[$, on a set Ω_δ verifying $\mathbb{P}_{\mathcal{L}_1}(\Omega_\delta) \geq 1 - \delta$,*

$$\mathbb{E} [l_{n_1}(f^*, \tilde{f}) \mid \mathcal{L}_1] \leq C' \inf_{T \leq T_{max}} \left\{ \inf_{f \in \mathcal{F}_T} l_{n_1}(f^*, f) + h^{-1} \log \left(\frac{n_1}{V} \right) \frac{|\tilde{T}|}{n_1} \right\} + h^{-1} \frac{C_\delta}{n_1} \quad (12)$$

$$+ h^{-1} C_2 \frac{\log K}{n_3}, \quad (13)$$

with $C_\delta = C'_1 + C''_1 \log(1/\delta)$.

Note that the constants appearing in the upper bounds for the risks are not sharp. We do not investigate the sharpness of the constants here.

Let us comment the results given in Theorem 1:

1) Both methods M1 and M2 are considered for the following reasons:

- Since all the risks are considered conditionally on the growing procedure, the M1 method permits to make a deterministic penalized model selection and then to obtain sharper upper bounds than the M2 method.
- On the other hand, the M2 method permits to keep the whole information given by \mathcal{L}_1 . Indeed, in that case, the sequence of pruned subtrees is not obtained via some plug-in method using a first split of the sample to provide the collection of tree-structured models. This method is the one proposed by Breiman *et al.* and it is more commonly applied in practice than the former. We focus on this method to ensure that it provides classifiers that have good performance in terms of risk.

2) For both the M1 and M2 methods, the inequality of Theorem 1 can be separated into two parts:

- (10) and (12) correspond to the pruning procedure. They show that, up to some absolute constant and the final selection, the conditional risk of the final classifier is approximately of the same order as the infimum of the penalized risks of the collection of subtrees of T_{max} . The term inside the infimum is of the same form as the penalized criterion (7) used in the pruning procedure. This shows that, for a sufficiently large temperature α , this criterion permits to select convenient subtrees in term of conditional risk. Let us emphasize that the penalty term is directly proportional to the number of leaves in the M1 method, whereas a multiplicative logarithmic term appears in the M2 method. This term is due to the randomness of the models considered, since the samples used to construct and prune T_{max} are no longer independent.
- (11) and (13) correspond to the final selection of \tilde{f} among the collection of pruned tree structured classifiers using \mathcal{L}_3 . As $K \leq n_1$, this selection adds a term proportional to $\log n_1/n_3$ for both methods, which shows that not much is lost when a test sample is used provided that n_3 is sufficiently large with respect to $\log n_1$. Nevertheless, since we have no idea of the size of the constant C_2 , it is difficult to deduce a general way of choosing \mathcal{L}_3 from this upper bound.

3) Let us comment the role of the Vapnik-Chervonenkis dimension of the set of splits \mathcal{S} used to construct T_{max} . Let us take the more often used case in CART, where \mathcal{S} is the set of all half-spaces of $\mathcal{X} = \mathbb{R}^d$. In this particular case, we have $V = d + 1$. So, if \mathcal{X} is low dimensional, the $\log n_1$ term has to be taken into account in the risk bound. Nevertheless, if CART provides models such that

- the maximal dimension of the models is $D_N = o(N/\log N)$,
- the approximation properties of the models are convenient enough to ensure that the bias tends to zero with increasing sample size N ,

then we have a result of consistency for \tilde{f} if n_3 is conveniently chosen with respect to $\log n_1$.

4) Let us emphasize the role of the margin in the quality of the selected classifier. Theorem 1 shows that the higher the margin, the smaller the risk, which is intuitive since the more separable the labels are, the easier the classification shall be. This confirms the fact that CART does a convenient job if margin assumption **MA(1a)** or **MA(1b)** is fulfilled.

Furthermore, let us comment briefly on the size of the margin to obtain oracle-type inequalities in Theorem 1. Massart *et al.* [29] show that, if $h \leq \sqrt{|\tilde{T}|/n}$ for one model \mathcal{F}_T (where $n = n_2$ for M1 and $n = n_1$ for M2), then the upper bound for the risk on this model (and then the penalty term in our framework) is of order $\sqrt{|\tilde{T}|/n}$. They obtain this result via minimax bounds for the risk that make a connection between the zero error case (corresponding to $h = 1/2$), with a minimax risk of order $|\tilde{T}|/n$, and the “global” pessimistic case (corresponding to $h = 0$), with a minimax risk of order $\sqrt{|\tilde{T}|/n}$.

These results suggest that Theorem 1 gives oracle-type inequalities only if $h > \sqrt{|\tilde{T}|/n}$ for every tree T pruned from T_{max} . Let us recall that the pruning procedure and consequently the results of Theorem 2.1.1 heavily depend on the linearity of the penalized criterion (7). It is not clear whether these results remain valid when using a non-linear penalty function, so we need to keep a penalty term of order $|\tilde{T}|/n$ to ensure that the sequence of pruned subtrees contains the whole statistical information. Hence CART will underpenalize trees for which $h \leq \sqrt{|\tilde{T}|/n}$, since in that case the penalty term should be of order $\sqrt{|\tilde{T}|/n} > |\tilde{T}|/n$. Due to the recursiveness of the pruning algorithm, if the above mentioned case occurs, then CART may select classifiers having an excessive number of leaves.

Nevertheless, the condition on the size of the margin can be forced via the growing procedure. Indeed, if the condition $h > \sqrt{|\tilde{T}_{max}|/n}$ is fulfilled, then the penalty is optimal in terms of risk. This condition can be controlled during the growing procedure by forcing the maximal tree’s construction to stop earlier for example. This is obviously difficult to do in practice since it heavily depends on the data and on the size of the learning sample, and is worth being investigated more deeply (on going work).

The two following subsections give more precise results on the pruning algorithm for both the M1 and M2 methods, and particularly on the constants appearing in the penalty function. Subsection 3.2 validates the discrete selection by test-sample. Note that the two results obtained for the validation of the pruning algorithm also hold in the case of deterministic X_i ’s.

3.1 Validation of the Pruning Procedure

In this section, we focus more particularly on the pruning algorithm and give trajectorial risk bounds for the classifier associated with T_α , the smallest minimizing subtree for the temperature α defined in subsection 3.1. We show that, for a convenient constant α , \hat{f}_{T_α} is not far from f^* in terms of its risk conditionally on \mathcal{L}_1 . Let us emphasize that the subsample \mathcal{L}_3 plays no role in the two following results.

3.1.1 \tilde{f} constructed via M1

Here we consider the second subsample \mathcal{L}_2 of n_2 observations. We assume that T_{max} is constructed on the first set of observations \mathcal{L}_1 and then pruned with the second set \mathcal{L}_2 independent of \mathcal{L}_1 . Since the set of pruned subtrees is deterministic according to \mathcal{L}_2 , we make a selection among a deterministic collection of models.

For any subtree T of T_{max} , let \mathcal{F}_T be the model defined on the leaves of T given by (8).

f^* will then be estimated on \mathcal{F}_T , whose dimension is $|\tilde{T}|$.

Then we choose the estimators as follows: let γ_{n_2} be the empirical contrast as defined by (5).

- For $T \preceq T_{max}$, $\hat{f}_T = \operatorname{argmin}_{f \in \mathcal{F}_T} [\gamma_{n_2}(f)]$,
- For $\alpha > 0$, T_α is the smallest minimizing subtree for the temperature α as defined in subsection 2.1.2 and $\hat{f}_{T_\alpha} = \operatorname{argmin}_{f \in \mathcal{F}_{T_\alpha}} [\gamma_{n_2}(f)]$.

Let us now consider the behaviour of \hat{f}_{T_α} .

Proposition 1. *Let $P_{\mathcal{L}_2}$ be the product distribution on \mathcal{L}_2 and let h be the margin given by MA(1a). Let $\xi > 0$.*

There exists a large enough positive constant $\alpha_0 > 2 + \log 2$ such that, if $\alpha > \alpha_0$, then there exist some nonnegative constants Σ_α and C such that

$$l(f^*, \hat{f}_{T_\alpha}) \leq C_1(\alpha) \inf_{T \preceq T_{max}} \left\{ \inf_{f \in \mathcal{F}_T} l(f^*, f) + h^{-1} \frac{|\tilde{T}|}{n_2} \right\} + C h^{-1} \frac{1 + \xi}{n_2}$$

on a set Ω_ξ such that $P_{\mathcal{L}_2}(\Omega_\xi) \geq 1 - \Sigma_\alpha e^{-\xi}$, where l is defined by (3), $C_1(\alpha) > \alpha_0$ and Σ_α are increasing with α .

We obtain a trajectorial non-asymptotic risk bound on a large probability set, leading to the conclusions given for Theorem 1. Nevertheless, taking an excessive temperature α will overpenalize and select a classifier having high risk $\mathbb{E}[l(f^*, \hat{f}_{T_\alpha}) \mid \mathcal{L}_1]$. Furthermore, the fact that $C_1(\alpha)$ and Σ_α are increasing with α suggests that both sides of the inequality grow with α . The choice of the convenient temperature is then critical to make a good compromise between the size of $\mathbb{E}[l(f^*, \hat{f}_{T_\alpha}) \mid \mathcal{L}_1]$ and a large enough penalty term.

In practice, since this temperature depends on the unknown margin h and some unknown constants, the use of a test sample as described in Section 1 is a convenient choice, as shown by Proposition 3.

3.1.2 \tilde{f} constructed via M2

In this subsection we define the different contrasts, expected loss and estimators exactly in the same way as in subsection 3.1.1, although l is replaced by the empirical expected loss on $X_1^{n_1} = \{X_i; (X_i, Y_i) \in \mathcal{L}_1\}$ defined by (9),

$$l_{n_1}(f^*, f) = \mathbb{E}_Y [\gamma_{n_1}(f) - \gamma_{n_1}(f^*)],$$

since the models and the evaluations of the empirical errors $\gamma_{n_1}(\hat{f}_T)$ are computed on the same grid $X_1^{n_1}$. In this case, we obtain nearly the same performance for \hat{f}_{T_α} despite the fact that the constant appearing in the penalty term can now depend on n_1 :

Proposition 2. *Let $P_{\mathcal{L}_1}$ be the product distribution on \mathcal{L}_1 , l_{n_1} (9) be the empirical expected loss computed on $\{X_i; (X_i, Y_i) \in \mathcal{L}_1\}$, and let h be the margin given by MA(1b). Let $\xi > 0$ and*

$$\alpha_{n_1, V} = 2 + V/2 \left(1 + \log \frac{n_1}{V} \right).$$

There exists a large enough positive constant α_0 such that, if $\alpha > \alpha_0$, then there exist some nonnegative constants Σ_α and C' such that

$$l_{n_1}(f^*, \hat{f}_{T_\alpha}) \leq C'_1(\alpha) \inf_{T \preceq T_{max}} \left\{ \inf_{f \in \mathcal{F}_T} l_{n_1}(f^*, f) + h^{-1} \alpha_{n_1, V} \frac{|\tilde{T}|}{n_1} \right\} + C' h^{-1} \frac{1 + \xi}{n_1}$$

on a set Ω_ξ such that $P_{\mathcal{L}_1}(\Omega_\xi) \geq 1 - 2\Sigma_\alpha e^{-\xi}$, where $C'_1(\alpha) > \alpha_0$ and Σ_α are increasing with α .

We obtain a similar trajectorial non-asymptotic risk bound on a large probability set. The same conclusions as those derived from the M1 case hold in this case. Let us just mention that the penalty term takes into account the complexity of the collection of trees having fixed number of leaves which can be constructed on $\{X_i; (X_i, Y_i) \in \mathcal{L}_1\}$. Since this complexity is controlled via the VC-dimension V , V necessarily appears in the penalty term. It differs from Proposition 1 in the sense that the models we consider are random, so this complexity has to be taken into account to obtain an uniform bound.

Example: Let us consider the case where \mathcal{S} is the set of all half-spaces of $\mathcal{X} = \mathbb{R}^d$ (which is the most common case in the CART algorithm). In this case, $V = d + 1$, consequently, if $n_1 > d + 1$, we obtain a penalty proportional to

$$\left(\frac{4 + (d + 1)(1 + \log[n_1/(d + 1)])}{2h} \right) \frac{|\tilde{T}|}{n_1}.$$

So, if CART provides some minimax estimator on a class of functions, the $\log n_1$ term always appears for f^* in this class when working in a linear space of low dimension.

As for the M1 case, since the temperature α depends on the unknown margin h and some unknown constants, the use of a test sample to select the final classifier among the sequence of pruned subtrees is a convenient choice, as shown by Proposition 3.

3.2 Final Selection

We focus here on the final step of the CART procedure: the selection of the classifier \tilde{f} among the collection of pruned subtrees given by the pruning procedure by using a test sample \mathcal{L}_3 . Given the sequence $(T_k)_{1 \leq k \leq K}$ pruned from T_{max} as defined in subsection 3.1, let us recall that \tilde{f} is defined by

$$\tilde{f} = \operatorname{argmin}_{\{\hat{f}_{T_k}; 1 \leq k \leq K\}} \left[\gamma_{n_3}(\hat{f}_{T_k}) \right].$$

The performance of this classifier can be compared to the performance of the collection of classifiers $(\hat{s}_{T_k})_{1 \leq k \leq K}$ by the following:

Proposition 3.

(i) if \tilde{f} is constructed via M1, let $\lambda = l$ and $R_{n_3}(f^*, \tilde{f}) = \mathbb{E} \left[\lambda(f^*, \tilde{f}) \mid \mathcal{L}_1, \mathcal{L}_2 \right]$.

(ii) if \tilde{f} is constructed via M2, let $\lambda = l_{n_1}$ and $R_{n_3}(f^*, \tilde{f}) = \mathbb{E} \left[\lambda(f^*, \tilde{f}) \mid \mathcal{L}_1 \right]$, where l_{n_1} is defined by (9).

For both cases, there exist three absolute constants $C'' > 1$, $C'_1 > 3/2$ and $C'_2 > 3/2$ such that

$$R_{n_3}(f^*, \tilde{f}) \leq C'' \inf_{1 \leq k \leq K} \lambda(f^*, \hat{f}_{T_k}) + C'_1 h^{-1} \frac{\log K}{n_3} + h^{-1} \frac{C'_2}{n_3},$$

where K is the number of pruned subtrees extracted during the pruning procedure.

4 Concluding Remarks

We have proven that CART provides convenient classifiers in terms of conditional risk under a margin condition. Nevertheless, as for the regression case, the properties of the growing procedure need to be analyzed to obtain full unconditional upper bounds.

The remarks made after Theorem 1 on the size of the margin h give some prospects for the application of CART in practice. These prospects may be for example

- using the slope heuristic (see for example [4] [3]) to select a classifier among a collection,
- searching for a robust manner to determine if the margin assumption is fulfilled, permitting to use the blind selection by test sample.

Some track to estimate the margin h if assumption **MA(1a)** or **MA(1b)** is fulfilled could be to use mixing procedures as boosting (see [8] [13] for example). Hence this estimate could be used in the penalized criterion to help find the convenient temperature. It could also give an idea of the difficulty to classify the considered data and henceforth to help choose the most adapted classification method.

5 Proofs

Let us start with a preliminary result.

5.1 Local Bound for Tree-Structured Classifiers

Let $(X, Y) \in \mathcal{X} \times \{0, 1\}$ be a pair of random variables and $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ be n independent copies of (X, Y) . Let $\|\cdot\|_n$ denote the empirical norm on $X_1^n = (X_i)_{1 \leq i \leq n}$. Then given two classifiers f and g , let us define

$$d_n^2(f, g) = \frac{1}{n} \sum_{i=1}^n (f(X_i) - g(X_i))^2 := \|f - g\|_n^2.$$

Let \mathcal{M}_n^* be the set of all possible tree-structured partitions that can be constructed on the grid X_1^n , corresponding to trees having all possible splits in \mathcal{S} and all possible forms without taking account of the response variable Y . So \mathcal{M}_n^* only depends on the grid X_1^n and is independent of the variables (Y_1, \dots, Y_n) . Hence, for a tree $T \in \mathcal{M}_n^*$, define

$$\mathcal{F}_T = \left\{ \sum_{t \in \tilde{T}} a_t \mathbb{1}_t ; (a_t) \in \{0, 1\}^{|\tilde{T}|} \right\},$$

where \tilde{T} refers the set of the leaves of T . Then, for any $f \in \mathcal{F}_T$ and any $\sigma > 0$, define

$$B_T(f, \sigma) = \{g \in \mathcal{F}_T ; d_n(f, g) \leq \sigma\}$$

For each classifier $f : \mathcal{X} \rightarrow \{0, 1\}$, let us define the empirical contrast of f recentered conditionally on X_1^n

$$\bar{\gamma}_n(f) = \gamma_n(f) - \mathbb{E}[\gamma_n(f) \mid X_1^n], \quad (14)$$

where γ_n is defined for any given classifier f by

$$\gamma_n(f) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{f(X_i) \neq Y_i}.$$

Remark 2. If γ_n is evaluated on a sample (X'_i) independent of X_1^n , it is easy to check that the bounds we obtain in what follows are still valid by defining the distance with respect to the marginal distribution of X instead of the empirical distribution.

We have the following result:

Lemma 1. For any $f \in \mathcal{F}_T$ and any $\sigma > 0$

$$\mathbb{E} \left[\sup_{g \in B_T(f, \sigma)} |\bar{\gamma}_n(g) - \bar{\gamma}_n(f)| \mid X_1^n \right] \leq 2 \sigma \sqrt{\frac{|\tilde{T}|}{n}}.$$

Proof. First of all, let us mention that, since the different variables we consider take values in $\{0; 1\}$, we have for all $x \in \mathcal{X}$ and all $y \in \{0, 1\}$

$$\mathbb{1}_{g(x) \neq y} - \mathbb{1}_{f(x) \neq y} = (g(x) - f(x))(1 - 2\mathbb{1}_{y=1}),$$

yielding

$$\bar{\gamma}_n(g) - \bar{\gamma}_n(f) = \frac{1}{n} \sum_{i=1}^n (g(X_i) - f(X_i)) (1 - 2\mathbb{1}_{Y_i=1}) - \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (g(X_i) - f(X_i)) (1 - 2\mathbb{1}_{Y_i=1}) \mid X_1^n \right].$$

Let us now consider a Rademacher sequence of random signs $(\varepsilon_i)_{1 \leq i \leq n}$ independent of $(X_i, Y_i)_{1 \leq i \leq n}$. Then one has by a symmetrization argument

$$\mathbb{E} \left[\sup_{g \in B_T(f, \sigma)} |\bar{\gamma}_n(g) - \bar{\gamma}_n(f)| \mid X_1^n \right] \leq \mathbb{E} \left[\sup_{g \in B_T(f, \sigma)} \frac{2}{n} \left| \sum_{i=1}^n \varepsilon_i (g(X_i) - f(X_i)) (1 - 2\mathbb{1}_{Y_i=1}) \right| \mid X_1^n \right].$$

Since g and f belong to \mathcal{F}_T , we have that

$$g - f = \sum_{t \in \tilde{T}} (a_t - b_t) \varphi_t,$$

where each (a_t, b_t) takes values in $[0, 1]^2$ and $(\varphi_t)_{t \in \tilde{T}}$ is an orthonormal basis of \mathcal{F}_T adapted to \tilde{T} (i.e. some normalized characteristic functions). Then by applying the Cauchy-Schwarz inequality, since $g \in B_T(f, \sigma)$, $\|g - f\|_n^2 = d_n^2(f, g) = \sum_{t \in \tilde{T}} (a_t - b_t)^2 \leq \sigma^2$, we obtain that

$$\begin{aligned} \left| \sum_{i=1}^n \varepsilon_i (g(X_i) - f(X_i)) (1 - 2\mathbb{1}_{Y_i=1}) \right| &\leq \sqrt{\sum_{t \in \tilde{T}} (a_t - b_t)^2} \sqrt{\sum_{t \in \tilde{T}} \left(\sum_{i=1}^n \varepsilon_i (1 - 2\mathbb{1}_{Y_i=1}) \varphi_t(X_i) \right)^2} \\ &\leq \sigma \sqrt{\sum_{t \in \tilde{T}} \left(\sum_{i=1}^n \varepsilon_i (1 - 2\mathbb{1}_{Y_i=1}) \varphi_t(X_i) \right)^2}. \end{aligned}$$

Finally, since $(\varepsilon_i)_{1 \leq i \leq n}$ and $(1 - 2\mathbb{1}_{Y_i=1})_{1 \leq i \leq n}$ take their values in $\{-1; 1\}$, $(\varepsilon_i)_{1 \leq i \leq n}$ are centered and independent of $(X_i, Y_i)_{1 \leq i \leq n}$, and since for each $t \in \tilde{T}$ $\|\varphi_t\|_n^2 = 1$, Jensen's inequality implies

$$\mathbb{E} \left[\sup_{g \in B_T(f, \sigma)} |\bar{\gamma}_n(g) - \bar{\gamma}_n(f)| \mid X_1^n \right] \leq 2 \frac{\sigma}{n} \sqrt{\sum_{t \in \tilde{T}} \sum_{i=1}^n \varphi_t^2(X_i)} \leq 2\sigma \sqrt{\frac{|\tilde{T}|}{n}}.$$

And the proof is achieved. \square

5.2 Proof of Proposition 1

To prove Proposition 1, we adapt results of Massart [27, Theorem 4.2], and Massart and Nédélec [29] (see also Massart *et.al.* [28]). Similar methods are used in [33].

Let $n = n_2$. Let us give a sample $\mathcal{L}_2 = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ of the random variable $(X, Y) \in \mathcal{X} \times [0, 1]$, where \mathcal{X} is a measurable space and let $f^* \in \mathcal{F} \subset \{f : \mathcal{X} \mapsto [0, 1] ; f \in \mathbb{L}^2(\mathcal{X})\}$ be the unknown function to be recovered. Assume $(\mathcal{F}_m)_{m \in \mathcal{M}_n}$ is a countable collection of countable models included in \mathcal{F} . Let us give a penalty function $\text{pen}_n : \mathcal{M}_n \rightarrow \mathbb{R}_+$, and $\gamma : \mathcal{F} \times (\mathcal{X} \times [0, 1]) \rightarrow \mathbb{R}_+$ a contrast function, i.e. γ such that $f \mapsto \mathbb{E}[\gamma(f, (X, Y))]$ is convex and minimum at point f^* . Hence define for all $f \in \mathcal{F}$ the expected loss $l(f^*, f) = \mathbb{E}[\gamma(f, (X, Y)) - \gamma(f^*, (X, Y))]$.

Finally let

$$\gamma_n = \frac{1}{n} \sum_{i=1}^n \gamma(\cdot, (X_i, Y_i)) \quad (15)$$

be the empirical contrast associated with γ . Let \hat{m} be defined as

$$\hat{m} = \operatorname{argmin}_{m \in \mathcal{M}_n} [\gamma_n(\hat{f}_m) + \text{pen}_n(m)]$$

where $\hat{f}_m = \operatorname{argmin}_{g \in \mathcal{F}_m} \gamma_n(g)$ is the minimum empirical contrast estimator of f^* on \mathcal{F}_m . Then the final estimator of f^* is

$$\tilde{f} = \hat{f}_{\hat{m}}. \quad (16)$$

One makes the following assumptions:

H₁: γ is bounded by 1, which is not a restriction since all the functions we consider take values in $[0, 1]$.

H₂: Assume there exist $c \geq (2\sqrt{2})^{-1/2}$ and some (pseudo-)distance d such that, for every pair $(f, g) \in \mathcal{F}^2$, one has

$$\operatorname{Var}[\gamma(g, (X, Y)) - \gamma(f, (X, Y))] \leq d^2(g, f),$$

and particularly for all $f \in \mathcal{F}$

$$d^2(f^*, f) \leq c^2 l(f^*, f).$$

H₃: For any positive σ and for any $f \in \mathcal{F}_m$, let us define

$$B_m(f, \sigma) = \{g \in \mathcal{F}_m ; d(f, g) \leq \sigma\}$$

where d is given by assumption **H₂**. Let $\bar{\gamma}_n = \gamma_n(\cdot) - \mathbb{E}[\gamma_n(\cdot)]$. We now assume that for any $m \in \mathcal{M}_n$, there exists some continuous function ϕ_m mapping \mathbb{R}_+ onto \mathbb{R}_+ such that $\phi_m(0) = 0$, $\phi_m(x)/x$ is non-increasing and

$$\mathbb{E} \left[\sup_{g \in B_m(f, \sigma)} |\bar{\gamma}_n(g) - \bar{\gamma}_n(f)| \right] \leq \phi_m(\sigma)$$

for every positive σ such that $\phi_m(\sigma) \leq \sigma^2$. Let ε_m be the unique solution of the equation $\phi_m(c\varepsilon) = \varepsilon^2$, $\varepsilon > 0$.

One gets the following result:

Theorem 2. Let $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ be a sample of independent realizations of the random pair $(X, Y) \in \mathcal{X} \times [0, 1]$. Let $(\mathcal{F}_m)_{m \in \mathcal{M}_n}$ be a countable collection of models included in some countable family $\mathcal{F} \subset \{f : \mathcal{X} \mapsto [0, 1] ; f \in \mathbb{L}^2(\mathcal{X})\}$. Consider some penalty function $\text{pen}_n : \mathcal{M}_n \rightarrow \mathbb{R}_+$ and the corresponding penalized estimator \tilde{f} (16) of the target function f^* . Take a family of weights $(x_m)_{m \in \mathcal{M}_n}$ such that

$$\Sigma = \sum_{m \in \mathcal{M}_n} e^{-x_m} < +\infty. \quad (17)$$

Assume that assumptions **H₁**, **H₂** and **H₃** hold.

Let $\xi > 0$. Hence, given some absolute constant $C > 1$, there exist some positive constants K_1 and K_2 such that, if for all $m \in \mathcal{M}_n$

$$\text{pen}_n(m) \geq K_1 \varepsilon_m^2 + K_2 c^2 \frac{x_m}{n},$$

then, with probability larger than $1 - \Sigma e^{-\xi}$,

$$l(f^*, \tilde{f}) \leq C \inf_{m \in \mathcal{M}_n} [l(f^*, \mathcal{F}_m) + \text{pen}_n(m)] + C' c^2 \frac{1 + \xi}{n},$$

where $l(f^*, \mathcal{F}_m) = \inf_{f_m \in \mathcal{F}_m} l(f^*, f_m)$ and the constant C' only depends on C .

Proof. The proof is inspired of Massart [27] and Massart *et.al.* [28]. We give only sketches of proofs since those are now routine results in the model selection area (see [28] for a fuller overview). The interested reader may find the detailed proofs in the first version of the paper [15].

Let $m \in \mathcal{M}_n$ and $f_m \in \mathcal{F}_m$. The definition of the expected loss and the fact that

$$\gamma_n(\tilde{f}) + \text{pen}_n(\hat{m}) \leq \gamma_n(f_m) + \text{pen}_n(m)$$

lead to the following inequality:

$$l(f^*, \tilde{f}) \leq l(f^*, f_m) + \bar{\gamma}_n(f_m) - \bar{\gamma}_n(\tilde{f}) + \text{pen}_n(m) - \text{pen}_n(\hat{m}) \quad (18)$$

where $\bar{\gamma}_n$ is defined by (14). The general principle is now to concentrate $\bar{\gamma}_n(f_m) - \bar{\gamma}_n(\tilde{f})$ around its expectation in order to offset the term $\text{pen}_n(\hat{m})$. Since $\hat{m} \in \mathcal{M}_n$, we proceed by bounding $\bar{\gamma}_n(f_m) - \bar{\gamma}_n(\hat{f}_{m'})$ uniformly in $m' \in \mathcal{M}_n$. For $m' \in \mathcal{M}_n$ and $f \in \mathcal{F}_{m'}$, let us define

$$w_{m'}(f) = \left[\sqrt{l(f^*, f_m)} + \sqrt{l(f^*, f)} \right]^2 + y_{m'}^2,$$

with $y_{m'} \geq \varepsilon_{m'}$, where $\varepsilon_{m'}$ is defined by assumption **H₃**. Hence let us define

$$V_{m'} = \sup_{f \in \mathcal{F}_{m'}} \frac{\bar{\gamma}_n(f_m) - \bar{\gamma}_n(f)}{w_{m'}(f)}.$$

Then (18) becomes

$$l(f^*, \tilde{f}) \leq l(f^*, f_m) + V_{\hat{m}} w_{\hat{m}}(\tilde{f}) + \text{pen}_n(m) - \text{pen}_n(\hat{m})$$

Since $V_{m'}$ can be written as

$$V_{m'} = \sup_{f \in \mathcal{F}_{m'}} \nu_n \left(\frac{\gamma(f_m, \cdot) - \gamma(f, \cdot)}{w_{m'}(f)} \right),$$

where ν_n is the recentered empirical measure, we bound $V_{m'}$ uniformly in $m' \in \mathcal{M}_n$ by using Rio's version of Talagrand's inequality recalled here: if \mathcal{F} is a countable family of measurable functions such that, for some positive constants v and b , one has for all $f \in \mathcal{F}$ $P(f^2) \leq v$ and $\|f\|_\infty \leq b$, then for every positive y , the following inequality holds for $Z = \sup_{f \in \mathcal{F}} (P_n - P)(f)$

$$\mathbb{P} \left[Z - \mathbb{E}(Z) \geq \sqrt{2 \frac{(v + 4b\mathbb{E}(Z))y}{n}} + \frac{by}{n} \right] \leq e^{-y}.$$

To proceed, we need to check the two bounding assumptions. First, since by assumption **H₁** the contrast γ is bounded by 1, we have that, for each $f \in \mathcal{F}_{m'}$,

$$\left| \frac{\gamma(f, \cdot) - \gamma(f_m, \cdot)}{w_{m'}(f)} \right| \leq \frac{1}{y_{m'}^2}. \quad (19)$$

Second, by using assumption **H₂**, we have that, for each $f \in \mathcal{F}_{m'}$,

$$\text{Var} \left[\frac{\gamma(f, (X, Y)) - \gamma(f_m, (X, Y))}{w_{m'}(f)} \right] \leq \frac{c^2}{4y_{m'}^2}. \quad (20)$$

Then, by Rio's inequality, we have for every $x > 0$

$$P \left[V_{m'} \geq \mathbb{E}(V_{m'}) + \sqrt{\frac{c^2 + 16\mathbb{E}(V_{m'})}{2ny_{m'}^2}} x + \frac{x}{ny_{m'}^2} \right] \leq e^{-x}.$$

Let us take $x = x_{m'} + \xi$, $\xi > 0$, where $x_{m'}$ is given by (17). Then by summing up over $m' \in \mathcal{M}_n$, we obtain that for all $m' \in \mathcal{M}_n$

$$V_{m'} \leq \mathbb{E}(V_{m'}) + \sqrt{\frac{c^2 + 16\mathbb{E}(V_{m'})}{2ny_{m'}^2}} (x_{m'} + \xi) + \frac{x_{m'} + \xi}{ny_{m'}^2}$$

on a set Ω_ξ such that $P(\Omega_\xi) \geq 1 - \Sigma e^{-\xi}$. We now need to bound $\mathbb{E}(V_{m'})$ in order to obtain an upper bound for $V_{m'}$ on the set of large probability Ω_ξ . By using techniques similar to Massart *et al.*'s [29], we obtain the following inequality via the monotonicity of $x \mapsto \phi(x)/x$ and the assumption $c \geq (2\sqrt{2})^{-1/2}$: for all $m' \in \mathcal{M}_n$,

$$\mathbb{E}[V_{m'}] \leq \frac{8\sqrt{10}\varepsilon_{m'} + c(2n)^{-1/2}}{y_{m'}}.$$

Hence, taking

$$y_{m'} = K \left[8\sqrt{10}\varepsilon_{m'} + c(2n)^{-1/2} + c\sqrt{\frac{x_{m'} + \xi}{n}} \right]$$

with $K > 0$, we obtain that, on Ω_ξ , for all $m' \in \mathcal{M}_n$,

$$V_{m'} \leq \frac{1}{K} \left[1 + \sqrt{\frac{1}{2} \left(1 + \frac{8}{K\sqrt{2}} \right)} + \frac{1}{2K\sqrt{2}} \right].$$

Finally, by using repeatedly the elementary inequality $(\alpha + \beta)^2 \leq 2\alpha^2 + 2\beta^2$ to bound $y_{\hat{m}}^2$ and $w_{\hat{m}}(f)$, we derive that the following inequality holds on Ω_ξ for any $m \in \mathcal{M}_n$ and any $f_m \in \mathcal{F}_m$:

$$\begin{aligned} (1 - 2K') l(f^*, \tilde{f}) &\leq (1 + 2K') l(f^*, f_m) + \text{pen}_n(m) + 2K'K^2 \frac{\xi}{n} + \frac{2c^2K'K^2}{n} \\ &\quad + 5 \times 2^9 K'K^2 \varepsilon_{\hat{m}}^2 + 2c^2K'K^2 \frac{x_{\hat{m}}}{n} - \text{pen}_n(\hat{m}), \end{aligned}$$

with

$$K' = \frac{C-1}{2(C+1)}, \quad K_1 = 5 \times 2^9 K' K^2, \quad K_2 = 2K' K^2,$$

achieving the proof. \square

Application to classification trees:

Let us now suppose that (X, Y) takes values in $\mathcal{X} \times \{0, 1\}$. The contrast is taken as $\gamma(f, (X, Y)) = \mathbb{1}_{f(X) \neq Y}$, the expected loss is defined by (3), and the collection of models is $(\mathcal{F}_T)_{T \leq T_{max}}$. The models and the collection are countable since there is a finite number of functions in each \mathcal{F}_T , and a finite number of nodes in T_{max} . Since we are working conditionally on \mathcal{L}_1 , we can apply Theorem 2 directly with \mathcal{L}_2 . To check assumption **H₂**, let us first note that, since all the variables we consider take values in $\{0, 1\}$, we have the following for all classifiers f and g

$$(\gamma(f, (X, Y)) - \gamma(g, (X, Y)))^2 = (\mathbb{1}_{Y \neq f(X)} - \mathbb{1}_{Y \neq g(X)})^2 \quad (21)$$

$$= (f(X) - g(X))^2. \quad (22)$$

Then if we take $d^2(f, g) = \mathbb{E}((f(X) - g(X))^2) = \|f - g\|^2$, where $\|\cdot\|$ is the \mathbb{L}^2 -norm with respect to the marginal distribution of X , we have that, for all classifiers f and g , $\text{Var}[\gamma(g, (X, Y)) - \gamma(f, (X, Y))] \leq d^2(f, g)$. Moreover, with the margin condition **MA(1a)**, we have that

$$l(f^*, f) \geq h \|f - f^*\|^2, \quad (23)$$

hence assumption **H₂** is checked with $d^2(f, g) = \|f - g\|^2$ and $c^2 = 1/h$, where h is the margin. By definition of h , we have $h \leq 1 \leq 2\sqrt{2}$, and then $c \geq (2\sqrt{2})^{-1/2}$.

Then assumption **H₃** is checked by Lemma 1 with $\phi_T(x) = 2x\sqrt{|\tilde{T}|/n}$. Hence Theorem 2 is verified with $\varepsilon_T = \sqrt{1/h}\sqrt{|\tilde{T}|/n}$.

Finally, to choose a convenient family of weights $(x_T)_{T \leq T_{max}}$, taking $x_T = \theta|\tilde{T}|$, with $\theta > 2 \log 2$ independent of $|\tilde{T}|$ as done in [17], we immediately obtain $\Sigma_\alpha = \Sigma_\theta < +\infty$. Then we get proposition 1 by Theorem 2.

5.3 Proof of Proposition 2

In what follows, we denote by \mathcal{L}_1 the sample $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ of size n of the random variable (X, Y) , and by X_1^n the sample $\{X_1, \dots, X_n\}$.

First we generalize Theorem 2 to random models, and then we apply it to CART. Let (X, Y) , \mathcal{F} $f^* \in \mathcal{F}$, $\mathcal{L}_1 = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$, γ and γ_n be defined as in subsection 5.2. Finally let us rewrite the expected loss of $f \in \mathcal{F}$ conditionally on X_1^n as

$$l_n(f^*, f) = \mathbb{E}_Y \left[\frac{1}{n} \sum_{i=1}^n (\gamma(f, (X_i, Y_i)) - \gamma(f^*, (X_i, Y_i))) \right],$$

where \mathbb{E}_Y is the expectation with respect to the marginal distribution of Y .

Let us consider a collection of at most countable models $(\mathcal{F}_m)_{m \in \mathcal{M}_n^*}$ and a subcollection $(\mathcal{F}_m)_{m \in \mathcal{M}_n}$, where $\mathcal{M}_n \subset \mathcal{M}_n^*$ may depend on $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$. Finally let us consider a penalty function $\text{pen}_n : \mathcal{M}_n \mapsto \mathbb{R}_+$ and let us define the estimator \hat{f} of f^* as follows: let

$$\hat{m} = \text{argmin}_{m \in \mathcal{M}_n} [\gamma_n(\hat{f}_m) + \text{pen}_n(m)],$$

where $\hat{f}_m = \operatorname{argmin}_{f \in \mathcal{F}_m} \gamma_n(f)$ is the minimum contrast estimator of f^* on \mathcal{F}_m . Then $\tilde{f} = \hat{f}_{\hat{m}}$.

Let us make the following assumptions.

H₁: γ is bounded by 1.

H₂: Assume there exist $c \geq (2\sqrt{2})^{-1/2}$ and some (pseudo-)distance d_n (that may depend on X_1^n) such that, for every pair $(g, f) \in \mathcal{F}$, one has

$$\operatorname{Var} [\gamma(g, (X, Y)) - \gamma(f, (X, Y)) \mid X_1^n] \leq d_n^2(g, f),$$

and particularly for all $f \in \mathcal{F}$

$$d_n^2(f^*, f) \leq c^2 l_n(f^*, f).$$

H₃: For any positive σ and for any $f \in \mathcal{F}_m$, let us define

$$B_m(f, \sigma) = \{g \in \mathcal{F}_m ; d_n(f, g) \leq \sigma\}$$

where d_n is given by assumption **H₂**. Let $\tilde{\gamma}_n$ be defined as (14). We now assume that for any $m \in \mathcal{M}_n$, there exists some continuous function ϕ_m mapping \mathbb{R}_+ onto \mathbb{R}_+ such that $\phi_m(0) = 0$, $\phi_m(x)/x$ is non-increasing and

$$\mathbb{E} \left[\sup_{g \in B_m(f, \sigma)} |\tilde{\gamma}_n(g) - \tilde{\gamma}_n(f)| \mid X_1^n \right] \leq \phi_m(\sigma)$$

for every positive σ such that $\phi_m(\sigma) \leq \sigma^2$. Let ε_m be the unique solution of the equation $\phi_m(c\varepsilon) = \varepsilon^2$, $\varepsilon > 0$.

One gets the following result.

Theorem 3. *Let $\mathcal{L}_1 = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ be a sample of independent realizations of the random pair $(X, Y) \in \mathcal{X} \times [0, 1]$. Let $(\mathcal{F}_m)_{m \in \mathcal{M}_n^*}$ be a countable collection of models included in some countable family $\mathcal{F} \subset \{f : \mathcal{X} \mapsto [0, 1] ; f \in \mathbb{L}^2(\mathcal{X})\}$ (which may depend on X_1^n). Consider some subcollection of models $(\mathcal{F}_m)_{m \in \mathcal{M}_n}$, where $\mathcal{M}_n \subset \mathcal{M}_n^*$ may depend on \mathcal{L}_1 , and some penalty function $\operatorname{pen}_n : \mathcal{M}_n \rightarrow \mathbb{R}_+$. Let \tilde{f} (16) be the corresponding penalized estimator of the target function f^* . Take a family of weights $(x_m)_{m \in \mathcal{M}_n^*}$ such that*

$$\sum_{m \in \mathcal{M}_n^*} e^{-x_m} \leq \Sigma < +\infty, \quad (24)$$

with Σ deterministic. Assume that assumptions **H₁**, **H₂** and **H₃** hold.

Let $\xi > 0$. Hence, given some absolute constant $C > 1$, there exist some positive constants K_1 and K_2 such that, if for all $m \in \mathcal{M}_n$

$$\operatorname{pen}_n(m) \geq K_1 \varepsilon_m^2 + K_2 c^2 \frac{x_m}{n},$$

then, with probability larger than $1 - 2\Sigma e^{-\xi}$,

$$l_n(f^*, \tilde{f}) \leq C \inf_{m \in \mathcal{M}_n} [l_n(f^*, \mathcal{F}_m) + \operatorname{pen}_n(m)] + C' c^2 \frac{1 + \xi}{n},$$

where $l_n(f^*, \mathcal{F}_m) = \inf_{f \in \mathcal{F}_m} l_n(f^*, f)$ and the constant C' only depends on C .

Proof. The proof is highly similar to that of Theorem 2. The main differences are in the conditioning and the fact that the collection of models $(\mathcal{F}_m)_{m \in \mathcal{M}_n}$ is random. To remove these issues, all the bounds are computed uniformly on \mathcal{M}_n^* so that the probability of the set we finally obtain is unconditional to X_1^n since Σ is deterministic. The inequalities are obtained by the same techniques as the ones used for the proof of the results on model selection on random models done by Gey and Nédélec in [17].

Let $m \in \mathcal{M}_n$ and $f_m \in \mathcal{F}_m$. Starting from (18), we have

$$l_n(f^*, \tilde{f}) \leq l_n(f^*, f_m) + w_{\hat{m}, m}(\tilde{f})V_{\hat{m}, m} + \text{pen}_n(m) - \text{pen}_n(\hat{m}), \quad (25)$$

where for all m' and M in \mathcal{M}_n^* , for all $f \in \mathcal{F}_{\uparrow m'}$ and $f_M \in \mathcal{F}_M$,

$$w_{m', M}(f) = \left[\sqrt{l(f^*, f)} + \sqrt{l_n(f^*, f_M)} \right]^2 + (y_{m'} + y_M)^2,$$

$$V_{m', M} = \sup_{f \in \mathcal{F}_{m'}} \left[\frac{\bar{\gamma}_n(f_M) - \bar{\gamma}_n(f)}{w_{m', M}(f)} \right],$$

with $y_{m'} \geq \varepsilon_{m'}$ and $y_M \geq \varepsilon_M$. The general principle is now exactly the same as in the proof of Theorem 2 despite the fact that we have to bound $V_{m', M}$ not only uniformly in $m' \in \mathcal{M}_n^*$, but also in $M \in \mathcal{M}_n^*$ in order to have an in-probability inequality that does not depend on X_1^n .

Assumption **H2** permits to give exactly the same upper bounds (except that they depend on X_1^n and that $y_{m'}$ is replaced by $y_{m'} + y_M$) as (19) and (20). By using the same techniques as in the proof of Theorem 2 and the same considerations as in [17], we obtain that

$$\begin{aligned} V_{m', M} &\leq \frac{1}{y_{m'} + y_M} \left(8\sqrt{10}\varepsilon_{m'} + \frac{c(2n)^{-1/2}}{2} + 8\sqrt{10}\varepsilon_M + \frac{c(2n)^{-1/2}}{2} \right) \\ &\quad + \sqrt{\frac{c^2 + 16(8\sqrt{10}(\varepsilon_{m'} + \varepsilon_M) + c(2n)^{-1/2})(y_{m'} + y_M)^{-1}}{2n(y_{m'}^2 + y_M^2)}} (x_{m'} + x_M + \xi) \\ &\quad + \frac{1}{y_{m'}^2 + y_M^2} \left(\frac{x_{m'} + \xi/2}{n} + \frac{x_M + \xi/2}{n} \right) \end{aligned}$$

on a set Ω_ξ such that $P(\Omega_\xi \mid X_1^n) \geq 1 - 2\Sigma e^{-\xi}$. Then, since Σ is deterministic, we get that $P(\Omega_\xi) \geq 1 - 2\Sigma e^{-\xi}$.

Hence, if we take for all $m' \in \mathcal{M}_n^*$

$$y_{m'} = 2K \left[8\sqrt{10}\varepsilon_{m'} + \frac{c(2n)^{-1/2}}{2} + c\sqrt{\frac{x_{m'} + \xi/2}{n}} \right],$$

we obtain that, on Ω_ξ , for all m' and M in \mathcal{M}_n^* ,

$$V_{m', M} \leq \frac{1}{K} \left[1 + \sqrt{\frac{1}{2} \left(1 + \frac{8}{K\sqrt{2}} \right)} + \frac{1}{2K\sqrt{2}} \right].$$

Finally the proof is achieved in the same way as the proof of Theorem 2. \square

Application to classification trees:

Let us consider the classification framework and the collection of models $(\mathcal{F}_T)_{T \leq T_{max}}$ obtained via the growing procedure in CART (see subsection 3.1) as recalled in subsection 5.2. Since the growing and the pruning procedures are made on the same sample \mathcal{L}_1 , the conditions of Theorem 3 hold. Since n_1 is fixed, let us consider \mathcal{M}_n^* as the set of all possible tree-structured partitions that can be constructed on the grid X_1^n , corresponding to trees having all possible splits in \mathcal{S} and all possible forms without taking account of the response variable Y . So \mathcal{M}_n^* depends only on the grid X_1^n and is independent of the variables (Y_1, \dots, Y_n) . Then $\{T \preceq T_{max}\} \subset \mathcal{M}_n^*$ and we are able to apply Theorem 3. Considering (21), we take $d_n(g, f) = \|f - g\|_n^2$, where $\|\cdot\|_n$ is the empirical norm on $X_1^{n_1}$. Using the margin condition **MA(1b)**, (23) is also verified for l_n and d_n , and we have assumption **H2** with $c^2 = 1/h$. Then, by Lemma 1, assumption **H3** is checked with $\phi_T(x) = 2x\sqrt{|\tilde{T}|/n}$ and, in the same way as in the proof of Proposition 1, ε_T is taken as $\varepsilon_T = \sqrt{1/h}\sqrt{|\tilde{T}|/n}$.

Finally, to choose a convenient family of weights $(x_T)_{T \in \mathcal{M}_n^*}$, taking (see [17])

$$x_T = V \left(\theta + \log \frac{n_1}{V} \right) |\tilde{T}|,$$

where V is the VC-dimension of the set of splits \mathcal{S} used to construct T_{max} and $\theta > 1$, we obtain

$$\Sigma_\alpha = \Sigma_\theta = \sum_{D \geq 1} \exp(-(\theta - 1)DV) < +\infty.$$

And we have Proposition 2.

5.4 Proof of Proposition 3

Proposition 3 is a direct application of the theorem obtained by Boucheron, Bousquet and Massart [7] recalled here: assume that we observe $N + n$ independent random variables with common distribution P depending on a parameter f^* to be estimated. Suppose the first N observations $Z' = Z'_1, \dots, Z'_N$ are used to build some preliminary collection of estimators $(\hat{f}_m)_{m \in \mathcal{M}_n}$ and the remaining observations Z_1, \dots, Z_n are used to select an estimator \tilde{f} among this collection by minimizing the empirical contrast as defined by (15) (with (X, Y) replaced by Z). Hence we have the following result.

Theorem 5.4.1 (Boucheron, Bousquet, Massart [7]).

Suppose that \mathcal{M}_n is finite with cardinal K . Assume that there exists some continuous function w mapping \mathbb{R}_+ onto \mathbb{R}_+ such that $x \mapsto w(x)/x$ is nonincreasing, and which satisfies for all $\varepsilon > 0$

$$\sup_{\{f \in \mathcal{F} ; l(f^*, f) \leq \varepsilon^2\}} \text{Var} [\gamma(f, Z) - \gamma(f^*, Z)] \leq w(\varepsilon). \quad (26)$$

Then one has for every $\theta \in (0, 1)$

$$(1 - \theta)\mathbb{E} [l(f^*, \tilde{f}) \mid Z'] \leq (1 + \theta) \inf_{m \in \mathcal{M}_n} l(f^*, \hat{f}_m) + \delta_*^2 \left(2\theta + (1 + \log(K)) \left(\frac{1}{3} + \frac{1}{\theta} \right) \right),$$

where l is defined by (3) and δ_* satisfies $\sqrt{n}\delta_*^2 = w(\delta_*)$.

Taking $w(\varepsilon) = (1/\sqrt{h})\varepsilon$ for both methods M1 and M2, where h is the margin, leads to proposition 3 with

$$C = \frac{1 + \theta}{1 - \theta}, \quad C_1 = \frac{\theta + 3}{2\theta(1 - \theta)}, \quad C_2 = C_1 + \frac{\theta}{1 - \theta}.$$

5.5 Proof of Theorem 1

We are now able to prove Theorem 1 via propositions 1, 2 and 3. The beginning of the proof remains the same if \tilde{f} is constructed either via M1 or M2. So we just give the first step of the proof for the M1 method.

Actually, since we have at most one model per dimension in the pruned subtree sequence, it suffices to note that $K \leq n_1$. Then let α_0 be the minimal constant given by Proposition 1. Hence, since for a given $\alpha > 0$ T_α belongs to the sequence $(T_k)_{1 \leq k \leq K}$,

$$\mathbb{E} \left[l(f^*, \tilde{f}) \mid \mathcal{L}_1, \mathcal{L}_2 \right] \leq C'' \inf_{\alpha > \alpha_0} l(f^*, \hat{f}_{T_\alpha}) + C'_1 h^{-1} \frac{\log K}{n_3} + h^{-1} \frac{C'_2}{n_3}.$$

Starting from this inequality, if \tilde{f} is constructed via M1, by using Proposition 1 with $\alpha = 2\alpha_0$ and by taking the expectation according to \mathcal{L}_2 , we obtain Theorem 1 with the appropriate constants.

Yet, if \tilde{f} is constructed via M2, we apply Proposition 2 with $\alpha = 2\alpha_0\alpha_{n_1, V}$ and, for each $\delta \in]0; 1[$, $\xi = \log(2\Sigma_\alpha/\delta)$. Then we obtain Theorem 1 with the appropriate constants.

References

- [1] AĪZERMAN, M. A., BRAVERMAN, E. M., AND ROZONOÈR, L. I. *Method of Potential Functions in the Theory of Learning Machines*. Nauka, Moscow (in Russian). 1970.
- [2] ARLOT, S., AND BARTLETT, P. Margin adaptive model selection in statistical learning. Tech. Rep. 0804.2937, arXiv, 2008.
- [3] ARLOT, S., AND MASSART, P. Data-driven calibration of penalties for least squares regression. Tech. Rep. 0802.0837, arXiv, 2008.
- [4] BIRGÉ, L., AND MASSART, P. A generalized cp criterion for gaussian model selection. Tech. Rep. 647, Universités de Paris 6 et 7, 2001.
- [5] BLANCHARD, G., SCHAFFER, C., ROZENHOLC, Y., AND MULLER, K.-R. Optimal dyadic decision trees. *Machine Learning* 66, 2-3 (2007), 209–242.
- [6] BOUCHERON, S., BOUSQUET, O., AND LUGOSI, G. Theory of classification: a survey of some recent advances. *ESAIM Probab. Stat.* 9 (2005), 323–375 (electronic).
- [7] BOUCHERON, S., BOUSQUET, O., AND MASSART, P. Data-driven penalties. Unpublished manuscript, 2004.
- [8] BREIMAN, L. Arcing classifiers. *Ann. Statist.* 26, 3 (1998), 801–849. With discussion and a rejoinder by the author.
- [9] BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A., AND STONE, C. J. *Classification And Regression Trees*. Chapman & Hall, 1984.
- [10] CHOU, P. A., LOOKABAUGH, T., AND GRAY, R. M. Optimal pruning with applications to tree-structured source coding and modeling. *IEEE Transactions on Information Theory* 35, 2 (1989), 299–315.

- [11] DEVROYE, L., GYÖRFI, L., AND LUGOSI, G. *A probabilistic theory of pattern recognition*, vol. 31 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1996.
- [12] DONOHO, D. L. CART and best-ortho-basis : A connection. *The Annals of Statistics* 25, 5 (1997), 1870–1911.
- [13] FREUND, Y., AND SCHAPIRE, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55, 1 (1997), 119–139.
- [14] GELFAND, S. B., RAVISHANKAR, C., AND DELP, E. J. An iterative growing and pruning algorithm for classification tree design. *IEEE Transactions on PAMI* 13, 2 (1991), 163–174.
- [15] GEY, S. Margin adaptive risk bounds for classification trees. Tech. Rep. 0902.3130, arXiv, 2009.
- [16] GEY, S., AND LEBARBIER, E. Using cart to detect multiple change-points in the mean for large samples. Tech. Rep. 12, SSB, 2008.
- [17] GEY, S., AND NEDELEC, E. Model selection for CART regression trees. *IEEE Trans. Inform. Theory* 51, 2 (2005), 658–670.
- [18] HASTIE, T., TIBSHIRANI, R., AND FRIEDMAN, J. *The elements of statistical learning*. Springer, 2001.
- [19] KOHLER, M., AND KRZYŻAK, A. On the rate of convergence of local averaging plug-in classification rules under a margin condition. *IEEE Trans. Inform. Theory* 53, 5 (2007), 1735–1742.
- [20] KOLTCHINSKII, V. Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Statist.* 34, 6 (2006), 2593–2656.
- [21] KOLTCHINSKII, V. Rejoinder: “Local Rademacher complexities and oracle inequalities in risk minimization” [Ann. Statist. 34 (2006), no. 6, 2593–2656]. *Ann. Statist.* 34, 6 (2006), 2697–2706.
- [22] LECUÉ, G. Simultaneous adaptation to the margin and to complexity in classification. *Ann. Statist.* 35, 4 (2007), 1698–1721.
- [23] LUGOSI, G. Pattern classification and learning theory. In *Principles of nonparametric learning (Udine, 2001)*, vol. 434 of *CISM Courses and Lectures*. Springer, Vienna, 2002, pp. 1–56.
- [24] LUGOSI, G., AND VAYATIS, N. On the Bayes-risk consistency of regularized boosting methods. *Ann. Statist.* 32, 1 (2004), 30–55.
- [25] MAMMEN, E., AND TSYBAKOV, A. B. Smooth discrimination analysis. *Ann. Statist.* 27, 6 (1999), 1808–1829.
- [26] MARY-HUARD, T. *Reduction de la Dimension et Selection de Modeles en Classification Supervisee*. PhD thesis, Universite de Paris-Sud, nb 8303, July 2006.

- [27] MASSART, P. Some applications of concentration inequalities to statistics. *Annales de la Faculté des Sciences de Toulouse* (2000).
- [28] MASSART, P. *Concentration inequalities and model selection*, vol. 1896 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard.
- [29] MASSART, P., AND NÉDÉLEC, É. Risk bounds for statistical learning. *Ann. Statist.* *34*, 5 (2006), 2326–2366.
- [30] NOBEL, A. B. Recursive partitioning to reduce distortion. *IEEE Trans. on Inform. Theory* *43*, 4 (1997), 1122–1133.
- [31] NOBEL, A. B. Analysis of a complexity-based pruning scheme for classification trees. *IEEE Trans. Inform. Theory* *48*, 8 (2002), 2362–2368.
- [32] NOBEL, A. B., AND OLSHEN, R. A. Termination and continuity of greedy growing for tree-structured vector quantizers. *IEEE Trans. on Inform. Theory* *42*, 1 (1996), 191–205.
- [33] SAUVÉ, M., AND TULEAU, C. Variable selection through cart. Tech. Rep. 5912, Institut National de Recherche en Informatique et en Automatique, 2006.
- [34] SCHAPIRE, R. E., FREUND, Y., BARTLETT, P., AND SUN LEE, W. Boosting the margin : a new explanation for the effectiveness of voting methods. *The Annals of Statistics* *26*, 5 (1998), 1651–1686.
- [35] TSYBAKOV, A. B. Optimal aggregation of classifiers in statistical learning. *Ann. Statist.* *32*, 1 (2004), 135–166.
- [36] TSYBAKOV, A. B., AND VAN DE GEER, S. A. Square root penalty: adaptation to the margin in classification and in edge estimation. *Ann. Statist.* *33*, 3 (2005), 1203–1224.
- [37] VAPNIK, V. N. *Statistical Learning Theory*. Wiley Inter-Sciences, 1998.
- [38] VAPNIK, V. N., AND CHERVONENKIS, A. Y. *Teoriya raspoznavaniya obrazov. Statisticheskie problemy obucheniya*. Izdat. “Nauka”, Moscow, 1974.
- [39] WERNECKE, POSSINGER, KALB, AND STEIN. Validating classification trees. *Biometrical Journal* *40*, 8 (1998), 993–1005.